arXiv:1405.5397v2 [math.PR] 7 Aug 2014

Anchored burning bijections on finite and infinite graphs

Samuel L. Gamlin* Antal A. Járai†

August 8, 2014

Abstract

Let G be an infinite graph such that each tree in the wired uniform spanning forest on G has one end almost surely. On such graphs G, we give a family of continuous, measure preserving, almost one-to-one mappings from the wired spanning forest on G to recurrent sandpiles on G, that we call anchored burning bijections. In the special case of $\mathbb{Z}^d$, $d \geq 2$, we show how the anchored bijection, combined with Wilson's stacks of arrows construction, as well as other known results on spanning trees, yields a power law upper bound on the rate of convergence to the sandpile measure along any exhaustion of $\mathbb{Z}^d$. We discuss some open problems related to these findings.

Key words: abelian sandpile, uniform spanning tree, loop-erased random walk, Wilson's algorithm, burning algorithm, wired spanning forest.

1 Introduction

The abelian sandpile model is a stochastic particle model defined on a graph by a cellular automaton. Also known as the chip-firing game [11], this model has shown interesting connections to a variety of different areas. In [3] the idea of self-organized criticality was introduced and the sandpile model was used as a simple example of the properties sought. Dhar [8] was the first to study the model in its own right, realising that many of its properties can be computed exactly, and hence it has the capacity to demonstrate important underlying principles of self-organized criticality. See the surveys [9, 30, 13] for background.

Let $G = (V \cup \{s\}, \mathsf{E})$ be a finite, connected multigraph, with a distinguished vertex s, called the sink. A *sandpile configuration*, that we usually denote by η, consists of assigning an integer number of particles $\eta(v) \in \{0, 1, 2, \dots\}$ to every non-sink vertex $v \in V$, The sandpile η is *stable*, if $\eta(v) \in \{0, 1, \dots, \deg_G(v) - 1\}$, for all $v \in V$, where $\deg_G(v)$ is the degree of the vertex v in the graph G.

*Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, United Kingdom, Email: S.Gamlin@bath.ac.uk

†Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, United Kingdom, Email: A.Jarai@bath.ac.uk

The dynamics of the model consist of two ingredients. The first is called *toppling*. This occurs when a vertex has at least as many particles as its degree. For such a vertex v, its height is reduced by its degree and one particle is sent along every edge incident with v (i.e. vertices with multiple edges connecting them to v receive more than one particle). Particles reaching the sink are lost (i.e. we do not keep track of them). The toppling of v is summarised by the mapping $\eta(w) \to \eta(w) - \Delta_{v,w}$, $w \in V$, using the Laplacian matrix

$$\Delta_{v,w} = \begin{cases} \deg_G(w) & \text{if } w = v; \\ -a_{v,w} & \text{if } w \neq v; \end{cases}$$

where $a_{v,w}$ is the number of edges connecting vertices $v, w \in V$. Starting from any sandpile η, carrying out all possible topplings (in any order) yields a unique stable sandpile η° [8].

The second ingredient of the model is *particle additions*. Given a stable sandpile η, we add a particle at a randomly chosen vertex $v \in V$, and then stabilize via topplings, if necessary. Successive particle additions yield a Markov chain on the set of stable sandpiles. We denote the set of recurrent states of this Markov chain by $\mathcal{R}_G$, and by ν_G the unique stationary distribution, that is the uniform distribution on $\mathcal{R}_G$ [8]. The following combinatorial characterization of $\mathcal{R}_G$ follows from [8, 25] (see also [11]):

$$\mathcal{R}_G = \left\{ \eta \in \prod_{x \in V} \{0, \ldots, \deg_G(x) - 1\} : \eta \text{ is ample for all } \emptyset \neq F \subset V \right\}.$$

Here η is called *ample for* F, if there exists $x \in F$ such that $\eta(x) \geq \deg_F(x)$.

An important tool for investigating sandpile configurations is the burning algorithm of Dhar [8], that we now describe. Given $\eta \in \mathcal{R}_G$, at time 0 we declare the sink to be "burnt". Following this, we successively "burn" vertices where $\eta(x)$ is at least as much as the number of edges leading from x to any unburnt neighbours. More precisely, we set

$$B_0 := \{s\}, \qquad U_0 := V,$$

and for $j \geq 1$ we inductively set:

$$B_j := \left\{ v \in U_{j-1} : \eta(v) \geq \deg_{U_{j-1}}(v) \right\}, \qquad U_j := U_{j-1} \setminus B_j.$$

Here B_j (resp. U_j), are the sets of vertices burnt, (resp. unburnt), at time j. Since η is ample for any non-empty U_{j-1}, we have $U_j = \emptyset$ eventually, at which time the algorithm terminates.

Majumdar and Dhar [25], following the above burning algorithm, constructed a bijection $\varphi_G : \mathcal{R}_G \to \mathcal{T}_G$, where $\mathcal{T}_G$ is the set of spanning trees of G. The map $\varphi_G : \eta \mapsto t$, that arises as a special case of the bijections introduced in Section 2, can be defined as follows. Fix for each $v \in V$ an ordering $\prec_v$ of the oriented edges $\{f : \text{tail}(f) = v\}$. If $v \in B_j$, let

$$\begin{aligned} m_v &:= \left| \left\{ f : \text{tail}(f) = v, \text{head}(f) \in \bigcup_{j' < j} B_{j'} \right\} \right|, \\ F_v &:= \{ e : \text{tail}(e) = v, \text{head}(e) \in B_{j-1} \}. \end{aligned} \tag{1}$$

Due to the burning rule, we have

$$\eta(v) = \deg(v) - m_v + \ell \quad \text{for some } 0 \le \ell < |F_v|. \tag{2}$$

With ℓ as above, let $e_v \in F_v$ be that edge e such that $|\{f \in F_v : f \prec_v e\}| = \ell$. Then we place, for each $v \in V$, the directed edge e_v in t, and forget the orientation. Observe that the burning time of a vertex $v \in V$ equals $\mathrm{dist}_t(v, s)$, where $\mathrm{dist}_t(\cdot, \cdot)$ is graph distance in t.

The image of ν_G under φ_G is the uniform spanning tree measure UST_G, i.e. the uniform distribution on $\mathcal{T}_G$. The burning bijection has been very fruitful in proving things about the sandpile model; see e.g. [29, 18, 2, 16]. It is natural to look for an extension of the burning bijection to infinite graphs, and this leads to some highly non-trivial questions. The main difficulty is that on finite graphs the burning algorithm starts from the sink, so the analogous process on infinite graphs should start from infinity. Our paper will be concerned with a particular way of overcoming this problem. However, as we outline below, some very natural questions remain open.

From now on, let $G = (V, \mathsf{E})$ be a locally finite, connected, infinite graph. Given an exhaustion by finite subgraphs: $V_1 \subset V_2 \subset \cdots \subset V$, $\cup_{n=1}^{\infty} V_n = V$, let $G_n = (V_n \cup \{s\}, \mathsf{E}_n)$ denote the *wired graph* obtained by identifying the vertices in $V \setminus V_n$, that becomes the sink s, and removing loop-edges at s. Note that there is a natural identification between E_n and those edges in E that have an endvertex in V_n. We denote by WSF the weak limit of the measures UST_{G_n} [23], called the *wired uniform spanning forest measure* on G. It is well-known, and easy to see, that WSF concentrates on spanning subgraphs of G all of whose components are infinite trees. We say that an infinite tree has *one end*, if any two infinite self-avoiding paths in the tree have a finite symmetric difference. We will assume that G satisfies the following condition:

$$\mathsf{WSF}\text{-a.s. all components have one end.} \tag{3}$$

While, in general, condition (3) is difficult to verify, it is known to hold on a large class of graphs, including $\mathbb{Z}^d$, $d \ge 2$; see [28, 7, 22, 23]. We denote

$$\mathcal{T} := \left\{ \begin{matrix} \text{spanning subgraphs of } G \text{ such that all} \\ \text{components are infinite one-ended trees} \end{matrix} \right\}.$$

The counterpart of $\mathcal{T}$ for the sandpile model will be

$$\mathcal{R} := \left\{ \eta \in \prod_{x \in V} \{0, \dots, \deg_G(x) - 1\} : \eta \text{ is ample for all finite } \emptyset \ne F \subset V \right\},$$

that we call the *recurrent configurations on* G.

Athreya and Járai [2] considered the case of $\mathbb{Z}^d$, $d \ge 2$, $V_n = [-n, n]^d \cap \mathbb{Z}^d$, and they showed that ν_{G_n} has a weak limit ν that concentrates on $\mathcal{R}$. When $2 \le d \le 4$ the argument is particularly transparent. It was shown by Pemantle [28] that when $2 \le d \le 4$, the measure WSF concentrates on the set

$$\mathcal{T}^{\mathrm{conn}} := \left\{ \text{one-ended spanning trees of } \mathbb{Z}^d \right\} \subset \mathcal{T}.$$

In this case the limiting sandpile measure ν is exhibited as the image of WSF under a map $\psi : \mathcal{T}^{\mathrm{conn}} \to \mathcal{R}$. Here ψ is defined essentially by inverting the relationships (1)–(2), that can be made sense of in $\mathbb{Z}^d$ for $t \in \mathcal{T}^{\mathrm{conn}}$. Namely, fix $t \in \mathcal{T}^{\mathrm{conn}}$ and $v \in \mathbb{Z}^d$. Let v^* denote the unique vertex such that all infinite paths starting at a neighbour of v pass through v^*, and v^* is nearest to v with respect to dist_t (such vertex exists because $t \in \mathcal{T}^{\mathrm{conn}}$). Orient all edges of t towards infinity (this makes sense, because t has one end). Let

$$m'_v := \left| \left\{ f : \mathrm{tail}(f) = v,\, \mathrm{dist}_t(\mathrm{head}(f), v^*) < \mathrm{dist}_t(v, v^*) \right\} \right|,$$
$$F'_v := \left\{ e : \mathrm{tail}(e) = v,\, \mathrm{dist}_t(\mathrm{head}(e), v^*) = \mathrm{dist}_t(v, v^*) - 1 \right\}.$$

Enumerate F'_v as $e_0 \prec_v \cdots \prec_v e_{|F'_v|-1}$, and let $0 \le \ell' < |F'_v|$ be the unique index such that $e_{\ell'} \in t$. Then we set

$$\psi(t)(v) := \eta(v) := 2d - m'_v + \ell', \quad v \in \mathbb{Z}^d.$$

It is not difficult to see that ψ is continuous on $\mathcal{T}^{\mathrm{conn}}$. (In a certain sense, ψ is the limit of the inverse bijections $\psi_{G_n} := \varphi_{G_n}^{-1} : \mathcal{T}_{G_n} \to \mathcal{R}_{G_n}$.) Moreover, ψ is equivariant under translations of $\mathbb{Z}^d$, if the orderings $\{\prec_v : v \in \mathbb{Z}^d\}$ are chosen equivariant. It is tempting to conjecture that ψ is almost one-to-one, i.e. injective up to sets of measure 0. We do not have a proof of this.

Open Question 1. Is ψ almost one-to-one in the case of $\mathbb{Z}^d$, $2 \le d \le 4$?

When $d > 4$, it turned out to be necessary to add extra randomness to the WSF in order to construct ν [2], so there is no natural mapping $\mathcal{T} \to \mathcal{R}$, a priori.

The main result of this paper is the construction of a family of measure preserving mappings between spanning forests and sandpiles that are almost one-to-one. Our mappings can be constructed on general infinite graphs G satisfying condition (3), in particular, also on some non-transitive graphs. In this general setting, Járai and Werning [15] showed that ν_{G_n} converges weakly to a limit ν, that is independent of the exhaustion. Our construction is a natural extension of the one in [15], that in turn was based on an observation of Majumdar and Dhar [24] and Priezzhev [29]. In general, when $G = (V, \mathsf{E})$ is transitive, our mappings will *not* be invariant under all graph automorphism.

Definition 1. An *anchor* is a sequence $\mathcal{D} = \{D_1, D_2, \ldots\}$ of finite subsets of vertices such that

(i) $D_1 \subset D_2 \subset \ldots$ and $\cup_{k \ge 1} D_k = V$;

(ii) D_k is simply connected for each $k \ge 1$, i.e. all connected components of $V \setminus D_k$ are infinite.

In Section 2 we will associate to any anchor $\mathcal{D}$ a particular burning rule. That is, for any finite $\Lambda \subset V$ and configuration $\eta \in \mathcal{R}_\Lambda$ we define burning times $\tau^{\mathcal{D}}(x, \Lambda; \eta)$, $x \in \Lambda$ in such a way that at each time only vertices that are burnable in the sense of Dhar [8] are burnt. The advantage of our rule will be that it is easy to pass to the limit $\Lambda \uparrow V$, i.e. we can define a consistent set of burning times $\tau^{\mathcal{D}}(x; \eta) \in \mathbb{Z}$ for ν-a.e. $\eta \in \mathcal{R}$. The reason for requiring (ii) in Definition 1 is that for general D_k, our burning rule will be identical if we replace D_k by the smallest simply connected set containing it.

Theorem 1. *Assume that the infinite graph G satisfies condition* (3). *The burning rule arising from any anchor $\mathcal{D}$ defines a continuous, measure preserving, injective map $\psi_{\mathcal{D}}$ from $(\mathcal{T}, \mathsf{WSF})$ to $(\mathcal{R}, \nu)$.*

The precise meaning of "defines" will become clear in Section 2, where we introduce the anchored bijection and the map $\psi_{\mathcal{D}}$. Indeed, the anchor will serve to prescribe a "preferred direction" for the burning of configurations on V starting from infinity. The following question complements Open Question 1.

Open Question 2. For $\mathbb{Z}^d$, $d > 4$, is there a continuous measure preserving map from $(\mathcal{T}, \mathsf{WSF})$ to $(\mathcal{R}, \nu)$ that is equivariant with respect to translations?

Open Questions 1 and 2 are connected to a result of Schmidt and Verbitskiy [31]. They constructed, for any $d \geq 2$, a family of $\mathbb{Z}^d$-equivariant continuous surjective mappings from $\mathcal{R}$ onto the so called harmonic model, i.e. functions from $\mathbb{Z}^d$ to the unit circle that are harmonic modulo 1. The image of ν under their maps is the unique measure of maximum entropy of the harmonic model [31, Theorem 5.9].

As an application of the anchored bijection, we show that combined with Wilson's stacks of arrows construction [32] it yields a coupling between ν_{G_n} and ν that we can analyze on $\mathbb{Z}^d$, $d \geq 2$. This leads to a power law upper bound on the rate of convergence of ν_{G_n} to ν.

Theorem 2. *Let $d \geq 2$ and let $\Lambda \subset \mathbb{Z}^d$ be finite. Let N be the radius of the largest ball centred at the origin that is contained in Λ. There exists $\alpha = \alpha(d) > 0$ such that for any $1 \leq k < N$ and any cylinder event E depending only on the heights within distance k of the origin we have*

$$|\nu_\Lambda(E) - \nu(E)| \leq C(k,d)N^{-\alpha}. \tag{4}$$

The exponent α and the dependence on k are explicit, although not optimal; see Theorem 9 and Theorem 17 for more detailed statements. Estimates analogous to (4), but restricted to $d = 2, 3$, have been given in the context of the zero dissipation limit in the abelian avalanche model [12, 17]. We believe that our approach will lead to a significant simplification, and an extension to all $d \geq 2$, of the arguments of [12].

As mentioned earlier, we will define burning processes on both finite and infinite configurations in such a way that these behave well with respect to taking limits. In particular, restricting an infinite recurrent configuration to distinct large finite sets Λ_1, Λ_2, the anchored burning processes on Λ_1 and Λ_2 couple with high probability, in the following sense:

$$\lim_{\Lambda \uparrow V} \left[\tau^{\mathcal{D}}(x, \Lambda; \eta) - \tau^{\mathcal{D}}(y, \Lambda; \eta)\right] = c(x, y; \eta). \tag{5}$$

We do not know whether the same statement is true for Dhar's original burning algorithm, where at each step *every* burnable vertex is burnt simultaneously.

Open Question 3. Let $\tau(x, k; \eta)$ denote the burning time of x with respect to Dhar's original burning algorithm in the ball of radius k centred at the origin in $\mathbb{Z}^d$. Does the analogue of (5) hold for $\mathbb{Z}^d$, $2 \leq d \leq 4$, as $k \to \infty$?

If the answer is yes, this would imply an affirmative answer to Open Question 1. This is because the coupling defines a burning time from infinity (unique up to a time shift) and this can be used to define the inverse map. Note that the arguments of [2] show that the statement of Open Question 3 fails for $\mathbb{Z}^d$, $d > 4$.

We close this introduction by remarking that a certain analogue of the statement of Open Question 3 holds on graphs of the form $G = G_0 \times \mathbb{Z}$, with G_0 a finite connected graph. Indeed,

with respect to the left-burnable measure studied by Járai and Lyons [14], it is not difficult to construct a configuration on a "triple of columns" $G_0 \times \{1,2,3\}$ that "synchronizes" burning from the left, and hence coupling occurs. It was in fact by studying this case that we arrived at the idea of anchored bijections.

The paper has the following structure. In Section 2 we define the anchored bijection in the finite case and then show how this extends to give a bijection in the infinite case. In Section 3 we present the quantitative bounds on $\mathbb{Z}^d$ when $d \geq 3$. In Section 4 we give the bounds on $\mathbb{Z}^2$. Throughout C_d will stand for an unspecified positive constant dependent only on d, that we change without any further indication.

2 Anchored bijections

Recall that $G = (V, \mathsf{E})$ is a locally finite infinite graph satisfying (3); we allow parallel edges. Let $\mathcal{D} = \{D_1, D_2, \dots\}$ be an anchor, and let $D_0 := \emptyset$. We call the set $E_k := D_k \setminus D_{k-1}$ the *k-th shell*. Given any finite $\Lambda \subset V$, we form the wired graph $G_\Lambda = (\Lambda \cup \{s\}, \mathsf{E}_\Lambda)$, and denote

$$\begin{aligned} \mathcal{T}_\Lambda &= \text{collection of spanning trees in } G_\Lambda, \\ \mathcal{R}_\Lambda &= \text{recurrent sandpiles in } G_\Lambda. \end{aligned}$$

We first define a bijection between $\mathcal{R}_\Lambda$ and $\mathcal{T}_\Lambda$ that is an extension of the one considered in [15].

Anchored bijection in finite Λ.

Let $K = \max\{k \geq 0 : D_k \subset \Lambda\}$. Fix $\eta \in \mathcal{R}_\Lambda$. Our definitions will depend on $\mathcal{D}$, but we will not always indicate this in our notation.

Phase 1. We apply the usual burning algorithm to η with the restriction that we do not allow any vertex of D_K to burn. That is, we define

$$\begin{aligned} B_0^{(1)} &:= \{s\}, \\ U_0^{(1)} &:= \Lambda, \end{aligned}$$

and for $j \geq 1$ we inductively set:

$$\begin{aligned} B_j^{(1)} &:= \left\{ v \in U_{j-1}^{(1)} \setminus D_K : \eta(v) \geq \deg_{U_{j-1}^{(1)}}(v) \right\}, \\ U_j^{(1)} &:= U_{j-1}^{(1)} \setminus B_j^{(1)}. \end{aligned}$$

We have $B_j^{(1)} = \emptyset$ eventually. Note that there may be vertices in $\Lambda \setminus D_K$ that do not burn in Phase 1. These vertices, together with the vertices in D_K, will burn in later phases.

Assuming Phase $i-1$ has already been defined for some $2 \leq i \leq K+1$, we inductively define Phase i as follows.

Phase i. We continue the burning algorithm on η with the restriction that no vertex of D_{K-i+1} is allowed to burn. That is, we set

$$B_0^{(i)} := \cup_{j\geq 0} B_j^{(i-1)},$$
$$U_0^{(i)} := \Lambda \setminus B_0^{(i)},$$

and for $j \geq 1$ we inductively set:

$$B_j^{(i)} := \left\{ v \in U_{j-1}^{(i)} \setminus D_{K-i+1} : \eta(v) \geq \deg_{U_{j-1}^{(i)}}(v) \right\},$$
$$U_j^{(i)} := U_{j-1}^{(i)} \setminus B_j^{(i)}.$$

We have $B_j^{(i)} = \emptyset$ eventually. Note that if $i \leq K$, there may be vertices in $\Lambda \setminus D_{K-i+1}$ that do not burn in Phase i, only later.

Since η is recurrent, all vertices that did not burn in Phases $1, \ldots, K$, do burn in Phase $K+1$ (if this was not true, we would have found a subset that is not ample for η). Hence we have $\cup_{j\geq 0} B_j^{(K+1)} = \Lambda \cup \{s\}$

We now define a map $\varphi_{\mathcal{D},\Lambda} : \mathcal{R}_\Lambda \to \mathcal{T}_\Lambda$. Regard G_Λ as an oriented graph, with each edge being present with both possible orientations. We fix for each $v \in \Lambda$ a linear ordering $\prec_v$ of the oriented edges e such that $\mathrm{tail}(e) = v$. Given the burning of η as above, we define what oriented edges will be present in the tree $t = \varphi_{\mathcal{D},\Lambda}(\eta)$.

If $v \in B_j^{(i)}$ for some $1 \leq i \leq K+1$ and $j \geq 1$, then we place an oriented edge pointing from v to some $w \in B_{j-1}^{(i)}$. In the case $j = 1$ such edge exists, because v must have a neighbour outside $U_0^{(i)}$, and hence in $B_0^{(i)}$. In the case $j \geq 2$ such edge also exists, because the requirement to burn v at step j implies that the degree of v in $U_{j-1}^{(i)}$ is strictly smaller than its degree in $U_{j-2}^{(i)}$. Hence v has a neighbour in $B_{j-1}^{(i)} = U_{j-2}^{(i)} \setminus U_{j-1}^{(i)}$. If there is more than one $w \in B_{j-1}^{(i)}$ neighbouring v, we make the choice of the edge dependent on $\eta(v)$, similarly to the usual burning bijection. Formally, we let:

$$m_v := \left| \left\{ f : \mathrm{tail}(f) = v,\, \mathrm{head}(f) \in \bigcup_{j'<j} B_{j'}^{(i)} \right\} \right|,$$
$$F_v := \left\{ e : \mathrm{tail}(e) = v,\, \mathrm{head}(e) \in B_{j-1}^{(i)} \right\}.$$

Due to the burning rule, we have

$$\eta(v) = \deg(v) - m_v + \ell \quad \text{for some } 0 \leq \ell < |F_v|.$$

With ℓ as above, let $e_v \in F_v$ be that edge e such that $|\{f \in F_v : f \prec_v e\}| = \ell$. Then we place the directed edge e_v in t.

Lemma 3. *For any $\eta \in \mathcal{R}_\Lambda$ the collection of edges t (disregarding their orientations) is a spanning tree of G_Λ, and the map $\varphi_{\mathcal{D},\Lambda} : \eta \mapsto t$ is injective. Consequently, $\varphi_{\mathcal{D},\Lambda}$ is a bijection between $\mathcal{R}_\Lambda$ and $\mathcal{T}_\Lambda$.*

Proof. It is clear from the definitions that there are no cycles in t, since the sets $B_j^{(i)}$, are disjoint and "lexicographically ordered" by the indices (i,j) for $1 \le i \le K+1$, $j \ge 1$. In order to show injectivity, suppose that $\eta_1 \neq \eta_2$. There is a first time (i,j) in the burning processes of η_1 and η_2, where the "two processes differ". That is, there exists a lexicographically smallest (i,j) such that $B_{j'}^{(i')}(\eta_1) = B_{j'}^{(i')}(\eta_2)$ for all $i' < i$, $j' \ge 1$ and for all $i' = i$, $j' < j$, and $\eta_1(v) = \eta_2(v)$ for all elements v of these sets, but

$$\text{there exists } v \in B_j^{(i)}(\eta_1) \cup B_j^{(i)}(\eta_2) \text{ such that } \eta_1(v) \neq \eta_2(v).$$

It is easy to check that our definition of $\varphi_{\mathcal{D},\Lambda}$ assigns different oriented edges emanating from v for η_1 and η_2. Since all edges are oriented towards the sink, this implies that the two trees also differ as unoriented trees, proving injectivity. Since $\mathcal{R}_\Lambda$ and $\mathcal{T}_\Lambda$ have the same number of elements $\det(\Delta)$ [8], it follows that $\varphi_{\mathcal{D},\Lambda}$ is a bijection. □

Given $\eta \in \mathcal{R}_\Lambda$, we define the *burning time* $\tau^{\mathcal{D}}(x,\Lambda;\eta)$ as the index of the pair (i,j) in the lexicographic order, where $B_j^{(i)} \ni x$, $1 \le i \le K+1$, $j \ge 1$ (we restrict to the non-empty $B_j^{(i)}$'s). Note that in general this differs from the graph distance of x from s in the tree $\varphi_{\mathcal{D},\Lambda}(\eta)$. This is because at Step 1 of Phase i, we may be connecting a vertex $v \in B_1^{(i)} \cap D_{K-i+2}$ to a vertex w that was burnt not in the last step of Phase $i-1$.

Given $D \subset \Lambda$ and a spanning tree t of G_Λ, we write $\mathrm{desc}_t(D)$ for the set of descendants of D in t, that is, the collection of vertices w such that the path in t from w to s has a vertex in D.

Lemma 4. *For any finite $\Lambda \subset V$, $1 \le i \le K+1$, and $\eta \in \mathcal{R}_\Lambda$, the set of vertices that did not burn by the end of Phase i are precisely the descendants of D_{K-i+1}. That is, we have $U_0^{(i+1)} = \mathrm{desc}_{\varphi_{\mathcal{D},\Lambda}(\eta)}(D_{K-i+1})$.*

Proof. Observe that all vertices in $B_1^{(i+1)}$ are in D_{K-i+1}, otherwise they could have been burnt in Phase i. Since the oriented edges assigned by the bijection respect the lexicographic order, and the orientation is towards the sink, this implies that all vertices burnt in Phases $i+1,\dots,K+1$ are in $\mathrm{desc}_{\varphi_{\mathcal{D},\Lambda}(\eta)}(D_{K-i+1})$. On the other hand, if a vertex v was burnt in one of the Phases $1,2,\dots,i$, then all vertices on the oriented path from v to s were also burnt in one of these Phases, and hence $v \notin \mathrm{desc}_{\varphi_{\mathcal{D},\Lambda}(\eta)}(D_{K-i+1})$. This completes the proof. □

We next formulate a consistency property between the sandpile configurations on the sets $\mathrm{desc}_{\varphi_{\mathcal{D},\Lambda}(\eta)}(D_k)$, $k \ge 1$, that will help us to take the limit $\Lambda \uparrow V$.

Definition 2. Given $k \ge 1$ and a finite simply connected set W with $D_k \subset W \subset V$, we define the graph $G^*_{W,k} = (W \cup \{s\}, \mathsf{E}^*_{W,k})$ as follows. It contains all the edges that W induces in the graph V, and for each edge $e \in \mathsf{E}$ that connects a vertex $u \in D_k$ with a vertex $v \in V \setminus W$, there is an edge in $\mathsf{E}^*_{W,k}$ between u and s. Note that there is a natural identification between $\mathsf{E}^*_{W,k}$ and a subset of E, and we will use this identification freely in what follows.

Lemma 5. *(i) Suppose $D_k \subset W \subset V$ with W simply connected. There is a mapping $\psi_{W,k} : \mathcal{T}_{G^*_{W,k}} \to \mathcal{R}_{G^*_{W,k}}$ such that whenever $\Lambda \supset W$, $t \in \mathcal{T}_\Lambda$ and $W = \mathrm{desc}_t(D_k)$ holds,*

$$\textit{the restriction of the sandpile } \varphi^{-1}_{\mathcal{D},\Lambda}(t) \textit{ to } W \textit{ equals } \psi_{W,k}(t_{W,k}),$$

where $t_{W,k}$ denotes the restriction of t to the edges in $\mathsf{E}^*_{W,k}$.
(ii) Suppose $D_{k'} \subset D_k \subset W$. Let $t \in \mathcal{T}_{G^*_{W,k}}$. If $W' = \mathrm{desc}_t(D_{k'})$, then

$$\textit{the restriction of } \psi_{W,k}(t) \textit{ to } W' \textit{ is given by } \psi_{W',k'}(t_{W',k'}).$$

Proof. (i) Write $\eta = \varphi^{-1}_{\mathcal{D},\Lambda}(t)$. Due to Lemma 4, the statement $W = \mathrm{desc}_t(D_k)$ is equivalent to the statement that in the sandpile η, W is precisely the set of vertices that did not burn in Phase $K - k + 1$. It is easy to check using the burning rules that as η varies over all sandpiles with this property, the restriction η_W ranges over $\mathcal{R}_{G^*_{W,k}}$, and $t_{W,k}$ is a spanning tree of $G^*_{W,k}$. It follows from our definition of Phases $K-k+2, \ldots, K+1$ of the anchored bijection that $t_{W,k}$ is entirely determined by η_W, in a way independent of Λ. The map $\eta_W \mapsto t_{W,k}$ is injective, and since $|\mathcal{R}_{G^*_{W,k}}| = |\mathcal{T}_{G^*_{W,k}}|$ it is bijective. Hence $\psi_{W,k}$ can be defined as the inverse of this map.

(ii) This follows similarly to part (i), because if $\Lambda \supset W$ and η is as in part (i), then the restriction of η_W to W' is $\eta_{W'}$. □

We are now ready to extend the bijection to G.

Anchored bijection on G.

Observe that for every $t \in \mathcal{T}$ and $v \in V$ there is a unique infinite path in t starting at v. Hence for any finite $D \subset V$, we can define $\mathrm{desc}_t(D)$ as those vertices for which the infinite path starting at v has a vertex in D.

Given $t \in \mathcal{T}$, for every $k \geq 1$ let $W_k = \mathrm{desc}_t(D_k)$. Observe that due to the one-end property (3) of elements of $\mathcal{T}$, W_k is finite for all $k \geq 1$. Denote by $t_{W_k,k}$ the restriction of t to the edges in $\mathsf{E}^*_{W_k,k}$. Due to Lemma 5(ii), the configurations $\psi_{W_k,k}(t_{W_k,k})$ consistently define a stable configuration η on V. This η will be an element of $\mathcal{R}$, because for any finite $F \subset V$ there exists $k \geq 1$ such that $D_k \supset F$, and $\psi_{W_k,k}(t_{W_k,k}) = \eta_{W_k}$ is ample for F. We denote the configuration obtained by $\psi_{\mathcal{D}}(t)$, so $\psi_{\mathcal{D}} : \mathcal{T} \to \mathcal{R}$.

Remark 1. Whenever $\Lambda \supset W_k = \mathrm{desc}_t(D_k)$, we have the following property. If we start burning $\psi_{\mathcal{D}}(t)|_\Lambda$ with the restriction that no vertex of D_k is allowed to burn, then the set of vertices that cannot be burnt is exactly W_k. This follows by considering the burning process in some $W_{k'} \supset \Lambda$.

Lemma 6. *The map $\psi_{\mathcal{D}}$ is injective and continuous.*

Proof. Suppose that $t_1, t_2 \in \mathcal{T}$ such that $\psi_{\mathcal{D}}(t_1) = \psi_{\mathcal{D}}(t_2)$. Let us denote $W_k^{(1)} = \mathrm{desc}_{t_1}(D_k)$ and $W_k^{(2)} = \mathrm{desc}_{t_2}(D_k)$, and let $\Lambda = W_k^{(1)} \cup W_k^{(2)}$. By Remark 1, if we start the burning process on $\psi_{\mathcal{D}}(t_1)|_\Lambda = \psi_{\mathcal{D}}(t_2)|_\Lambda$ in Λ (with the restriction that D_k is not allowed to burn), then the set of vertices that do not burn equals both $W_k^{(1)}$ and $W_k^{(2)}$. In particular, these sets are equal, that is, $W_k^{(1)} = W_k^{(2)}$. Denoting their common value by W_k, we have

$$\psi_{W_k,k}(t_1|_{\mathsf{E}^*_{W_k,k}}) = \psi_{\mathcal{D}}(t_1)|_{W_k} = \psi_{\mathcal{D}}(t_2)|_{W_k} = \psi_{W_k,k}(t_2|_{\mathsf{E}^*_{W_k,k}}).$$

Hence t_1 equals t_2 on $\mathsf{E}^*_{W_k,k}$. Since k is arbitrary, it follows that $t_1 = t_2$, and therefore $\psi_{\mathcal{D}}$ is injective.

In order to see continuity, fix $t \in \mathcal{T}$, let $\eta = \psi_{\mathcal{D}}(t)$, and let $k \geq 1$ be fixed. Let $W_k = \mathrm{desc}_t(D_k)$. Suppose that $t' \in \mathcal{T}$ has the property that t' agrees with t on all edges in E that have an end vertex in W_k. Then it follows that $\mathrm{desc}_{t'}(D_k) = W_k$, and $t'_{W_k,k} = t_{W_k,k}$. Therefore

$$\psi_{\mathcal{D}}(t')|_{W_k} = \psi_{W_k,k}(t'_{W_k,k}) = \psi_{W_k,k}(t_{W_k,k}) = \psi_{\mathcal{D}}(t)|_{W_k}.$$

Since $k \geq 1$ is arbitrary, $W_k \supset D_k$ and $\cup_{k\geq 1} D_k = V$, this implies continuity of $\psi_{\mathcal{D}}$. □

The following lemma follows directly from the proof of [15, Theorem 3]. We provide a sketch of the proof for the reader's convenience.

Lemma 7. *The image of* WSF *under* $\psi_{\mathcal{D}}$ *equals* $\nu = \lim_{\Lambda\uparrow V} \nu_\Lambda$.

Sketch of the proof. Let E be a cylinder event that only depends on the sandpile heights in D_k for some $k \geq 1$. For any $\Lambda \supset D_k$, let $W_{\Lambda,k}$ be the random set of vertices that are unburnt just before the phase in which we first allow vertices in D_k to burn, that is, $U_0^{(K-k+2)}$. Due to Lemma 4, $W_{\Lambda,k}$ also equals the set of descendants of D_k in $\psi^{-1}_{\mathcal{D},\Lambda}(\eta_\Lambda)$, where η_Λ is the sandpile configuration in Λ. Recall the auxiulliary graph $G^*_{W,k}$ from Definition 2. Due to the proof of Lemma 5(i), for any fixed set $D_k \subset W \subset \Lambda$, the conditional distribution of η_W, given the event $\{W_{\Lambda,k} = W\}$ is given by $\nu_{G^*_{W,k}}$. Hence, conditioning on the value of $W_{\Lambda,k}$, we have:

$$\nu_\Lambda(E) = \sum_{D_k \subset W \subset \Lambda} \nu_\Lambda(W_{\Lambda,k} = W)\nu_{G^*_{W,k}}(\eta_W \in E). \tag{6}$$

Note that, in the notation of Lemma 5, we have

$$\begin{aligned}\nu_{G^*_{W,k}}(\eta_W \in E) &= \mathsf{UST}_{G^*_{W,k}}(t : \psi_{W,k}(t) \in E)\\ &= \mathsf{WSF}(t : \psi_{W,k}(t_{W,k}) \in E \,|\, \mathrm{desc}_t(D_k) = W)\\ &= \mathsf{WSF}(t : \psi_{\mathcal{D}}(t) \in E \,|\, \mathrm{desc}_t(D_k) = W).\end{aligned}$$

In particular, this probability does not depend on Λ. We also have

$$\lim_{\Lambda\uparrow V} \nu_\Lambda(W_{\Lambda,k} = W) = \lim_{\Lambda\uparrow V} \mathsf{UST}_\Lambda(t : \mathrm{desc}_t(D_k) = W) = \mathsf{WSF}(t : \mathrm{desc}_t(D_k) = W).$$

This is because for a fixed finite set W, the event $\mathrm{desc}_t(D_k) = W$ is spanning-tree-local: it only depends on the status of the edges in $\mathsf{E}^*_{W,k}$. Finally, note that due to the one-end property (3) the family $\{W_{\Lambda,k} : \Lambda \supset D_k\}$ is tight, in the sense that

$$\lim_{M\to\infty} \sup_{\Lambda \supset D_M} \mathsf{UST}_\Lambda(t : \mathrm{desc}_t(D_k) \not\subset D_M) = 0.$$

This allows us to pass to the limit in (6) and obtain

$$\begin{aligned}\lim_{\Lambda\uparrow V} \nu_\Lambda(E) &= \nu(E)\\ &= \sum_{\substack{W : W \text{ is finite}\\ W \supset D_k}} \mathsf{WSF}(t : \mathrm{desc}_t(D_k) = W)\, \mathsf{WSF}(\psi_{\mathcal{D}}(t) \in E \,|\, \mathrm{desc}_t(D_k) = W)\\ &= \mathsf{WSF}(t : \psi_{\mathcal{D}}(t) \in E).\end{aligned}$$

□

Lemmas 6, 7 imply Theorem 1.

Our final lemma shows the coupling property (5).

Lemma 8. *Fix $o \in D_1$. For any $t \in \mathcal{T}$ and $x \in V$ the limit*

$$\lim_{\Lambda \uparrow V} \left[\tau^{\mathcal{D}}(x, \Lambda; \psi_{\mathcal{D}}(t)) - \tau^{\mathcal{D}}(o, \Lambda; \psi_{\mathcal{D}}(t))\right] =: \tau^{\mathcal{D}}(x; \eta) \in \mathbb{Z}$$

exist.

Proof. Let $k \geq 1$ be the smallest index such that $x \in D_k$, let $W = \mathrm{desc}_t(D_k)$, and suppose that $\Lambda \supset W$. Due to Remark 1, for any such Λ the last $k+1$ phases of the burning of η_Λ have identical history. This implies the claim. □

3 Rate of convergence in $\mathbb{Z}^d$, $d \geq 3$.

Henceforth we consider the graphs $G = \mathbb{Z}^d$, and in this section we assume $d \geq 3$. Let D_k be the intersection of the Euclidean ball of radius k about the origin with $\mathbb{Z}^d$. We write ∂W for the set of vertices in $W^c := \mathbb{Z}^d \setminus W$ that have a neighbour in W.

Let $\Lambda \subset \mathbb{Z}^d$ be finite. We consider the realizations of WSF and UST_{G_Λ} via stacks of arrows, as introduced by Wilson [32]; see also [23]. To each vertex $v \in \mathbb{Z}^d$ we associate an i.i.d. sequence of arrows $\{e_i^v : i = 1, 2, \dots\}$, where e_i^v is an oriented edge with $\mathrm{tail}(e_i^v) = v$ and $\mathrm{head}(e_i^v)$ uniformly random among the neighbours of v. The stacks associated to different v are independent. We define $\mathbf{P}$ as the underlying probability measure for the stacks of arrows. We say that e_i^v has colour i, and we envision e_1^v lying directly above e_2^v in the stack, and similarly, for all k, e_k^v lying directly above e_{k+1}^v. An oriented cycle $\mathcal{C}$ in $\mathbb{Z}^d$ is associated the *weight* $w(\mathcal{C}) = (2d)^{-|\mathcal{C}|}$, where $|\mathcal{C}|$ denotes the number of arrows in $\mathcal{C}$. Sometimes we will need to consider *coloured cycles*, that is, a cycle consisting of some arrows $e_{i_1}^{v_1}, \dots, e_{i_r}^{v_r}$. We will use bold characters, like $\mathbf{C}$, to denote coloured cycles. In this case, $\mathcal{C}$ will denote the cycle obtained from $\mathbf{C}$ by ignoring the colours.

Wilson's algorithm [32] is based on the idea of *cycle popping* that we now decribe. We start with a configuration of stacks of arrows, as described above. We say that initially e_i^v is *in position i*. We refer to the arrows in position 1 as lying on top of the stack. Suppose that arrows $e_1^{v_1}, \dots, e_1^{v_r}$ on top of the stacks form a coloured cycle $\mathbf{C}$. By *popping* $\mathbf{C}$, we mean removing the arrows in $\mathbf{C}$ from the stacks, and shifting the positions of the arrows beneath them upwards. That is: after popping $\mathbf{C}$, $e_k^{v_j}$ will be in position $k-1$ for $j = 1, \dots, r$, $k \geq 2$. Similarly, if at any later time some arrows $e_{i_1}^{v_1}, \dots, e_{i_r}^{v_r}$ are all in position 1 and form an oriented cycle $\mathbf{C}$, we may pop them and shift the arrows beneath them upwards.

As shown in [32], with probability 1, only finitely many coloured cycles contained in Λ can be popped, and on this event, regardless of what order of popping is chosen, the same set of coloured cycles get popped. Moreover, the arrows that are left on top of the stacks when no more cycles can be popped form a spanning tree of G_Λ (oriented towards s), that also does not depend on the order of popping. Furthermore, the set of coloured cycles popped and the tree obtained are independent, and the tree is distributed according to UST_{G_Λ}.

Cycle popping can also be made sense of in $\mathbb{Z}^d$, $d \geq 3$. One way is to use loop-erased random walks (LERW), as in [7, Theorem 5.1], known as Wilson's method rooted at infinity. Given a

finite path $\pi = [x_0, \dots, x_M]$ in $\mathbb{Z}^d$, its *loop-erasure* $\mathrm{LE}(\pi)$ is defined by chronologically erasing cycles from the path, as they are created; see [23]. Loop-erasure also makes sense for infinite paths π, as long as π visits every vertex finitely often. To describe Wilson's method rooted at infinity, order the vertices of $\mathbb{Z}^d$ arbitrarily as $v_1, v_2, \dots$. Starting from v_1, follow the arrows on top of the stacks, and whenever a cycle is completed, pop that cycle. The trajectory traced by this walk is a simple random walk $\{S^{(1)}(m)\}_{m\geq 0}$ under $\mathbf{P}$, so due to transience, every vertex is visited only finitely many times, with probability 1. Hence, on this event, there is a well-defined configuration of stacks of un-popped arrows, after the entire trajectory of $S^{(1)}$ has been traced. On top of the stacks now lie $\mathfrak{F}_1 := \mathrm{LE}(S^{(1)}[0,\infty))$, and unexamined arrows everywhere else. Next, starting from v_2, again follow the arrows on top of the stacks, popping any cycles that are completed. The trace of the path will now be a random walk $S^{(2)}[0,\tau^{(2)}]$, where $\tau^{(2)} \in [0,\infty]$ is the first hitting time of $\mathfrak{F}_1$. Upon hitting $\mathfrak{F}_1$, a segment of $\mathfrak{F}_1$ is retraced without encountering any further cycle, and on top of the stacks will lie $\mathfrak{F}_2 := \mathfrak{F}_1 \cup \mathrm{LE}(S^{(2)}[0,\tau^{(2)}])$, with unexamined arrows everywhere else. Continue this way with $v_3, v_4, \dots$. With probabilty one, from each stack only finitely many arrows get popped, hence the procedure reveals a random spanning forest T. Due to [7, Theorem 5.1], T is distributed according to WSF. We will also need the following alternative way of popping cycles in $\mathbb{Z}^d$:

$$\text{first pop all cycles contained in } D_1\text{, then pop all cycles contained in } D_2\text{, etc.} \tag{7}$$

Wilson's proof for finite graphs [32] can be adapted to show that on the probability 1 event when T is well-defined, the procedure (7) reveals exactly the same forest T. In particular, for any finite $\Lambda \subset \mathbb{Z}^d$, cycle popping in Λ also terminates with probability 1, resulting in a spanning tree T_Λ, distributed according to UST_{G_Λ}. Thus, using the same stacks of arrows for cycle popping in Λ and in $\mathbb{Z}^d$ provides the required coupling of WSF and UST_{G_Λ}.

Given a cylinder event $E \subset \{0,\dots,2d-1\}^{D_k}$ only depending on sandpile heights in D_k, let us write $E_{\mathbb{Z}^d} = \{\psi_{\mathcal{D}}(T) \in E\}$ and $E_\Lambda = \{\psi_{\mathcal{D},\Lambda}(T_\Lambda) \in E\}$. We have $\mathbf{P}(E_\Lambda) = \nu_\Lambda(E)$, due to Lemma 3 and $\mathbf{P}(E_{\mathbb{Z}^d}) = \nu(E)$, due to Lemma 7.

Theorem 9. *Let E be a cyclinder event depending only on the sandpile heights in D_k. Let $d \geq 3$, let $\Lambda \subset \mathbb{Z}^d$ be a finite set and let N be the radius of the largest ball centered at the origin that is contained in Λ. We have*

$$|\nu_\Lambda(E) - \nu(E)| \leq \mathbf{P}(E_\Lambda \Delta E_{\mathbb{Z}^d}) \leq \begin{cases} C_d k^{d-1} N^{\frac{2-d}{2d}} & \text{if } d \geq 5; \\ C k^{26/9} N^{-2/9} & \text{if } d = 4; \\ C k^{17/9} N^{-1/9} & \text{if } d = 3. \end{cases}$$

Here Δ denotes symmetric difference.

The proof is broken down into a number of propositions and lemmas. Let us write W_k for the random set of descendants of D_k in T,

Proposition 10. *Suppose $d \geq 3$, $1 \leq k < n < N$, and $\Lambda \supset D_N$. There is a constant $C_d > 0$ such that*

$$\mathbf{P}\left(W_k \subset D_n \text{ but } W_k \neq W_{k,\Lambda} \text{ or } T|_{\mathsf{E}^*_{W,k}} \neq T_\Lambda|_{\mathsf{E}^*_{W,k}}\right) \leq C_d \frac{k^{d-2} n^2}{(N-n)^{d-2}}. \tag{8}$$

Proof. If we successively pop all cycles in D_n, then in D_{n+1}, then in D_{n+2}, etc., then we see that $\mathbf{P}$-a.s. on the event $W_k \subset D_n$ we have $W_{k,\Lambda'} = W_k$ and $T|_{E^*_{W,k}} = T_\Lambda|_{E^*_{W,k}}$ for all large enough finite Λ'. Therefore, it is enough to show that for all finite $\Lambda' \supset \Lambda$ we have

$$\mathbf{P}\left(W_{k,\Lambda'} \subset D_n \text{ but } W_{k,\Lambda'} \neq W_{k,\Lambda} \text{ or } T_{\Lambda'}|_{\mathsf{E}^*_{W,k}} \neq T_\Lambda|_{\mathsf{E}^*_{W,k}}\right) \leq C_d \frac{k^{d-2}n^2}{N^{d-2}}, \tag{9}$$

with C_d independent of Λ, Λ'.

In order to prove (9), we first pop all cycles we can that are contained in Λ. This leaves on top of the stacks in Λ the wired spanning tree T_Λ of G_Λ. Let $\mathfrak{L}$ denote the collection of remaining coloured cycles contained in Λ' that need to be popped in order to obtain the wired spanning tree $T_{\Lambda'}$ in Λ'. For convenience, the cycles in $\mathfrak{L}$ are regarded as having colours according to their *current positions* in the stacks, i.e. after all cycles contained in Λ have been popped. We claim that the probability distribution of $\mathfrak{L}$ is proportional to total weight and that $\mathfrak{L}$ is independent of the wired spanning tree $T_{\Lambda'}$ in Λ', that is:

$$\mathbf{P}(\mathfrak{L} = \{\mathbf{C}_1, \ldots, \mathbf{C}_K\},\, T_{\Lambda'} = t_{\Lambda'}) = \mathsf{UST}_{G_{\Lambda'}}(t_{\Lambda'}) \frac{1}{Z} \prod_{j=1}^{K} w(\mathcal{C}_j), \tag{10}$$

where Z is a normalization factor. Indeed, we show that this follows from Wilson's theorem [32]. Let us write $\mathfrak{L}^0_\Lambda$, respectively $\mathfrak{L}^0_{\Lambda'}$, for the collection of coloured cycles contained in Λ, respectively Λ', that we need to pop in order to reveal T_Λ, respectively $T_{\Lambda'}$. Then $\mathfrak{L}$ is a deterministic function of $\mathfrak{L}^0_{\Lambda'}$ (recall that the colours of cycles in $\mathfrak{L}$ are according to their positions aquired after cycle popping in Λ is complete). By Wilson's theorem, $T_{\Lambda'}$ is independent of $\mathfrak{L}^0_{\Lambda'}$, and hence of $\mathfrak{L}$, and is distributed according to $\mathsf{UST}_{G_{\Lambda'}}$. Therefore, the left hand side of (10) equals

$$\mathsf{UST}_{G_{\Lambda'}}(t_{\Lambda'})\, \mathbf{P}(\mathfrak{L} = \{\mathbf{C}_1, \ldots, \mathbf{C}_K\}).$$

In order to show that the second factor is proportional to weight, first observe that $\mathfrak{L}^0_{\Lambda'}$ and the pair $(\mathfrak{L}^0_\Lambda, \mathfrak{L})$ are deterministic functions of each other. We show that $\mathfrak{L}^0_\Lambda$ and $\mathfrak{L}$ are independent. This is because, using Wilson's Theorem again, $\mathfrak{L}^0_\Lambda$, T_Λ, the stacks of arrows beneath T_Λ, and the stacks of arrows in $\Lambda' \setminus \Lambda$ are mutually independent, and $\mathfrak{L}$ is a deterministic function of the latter three. We have

$$\mathbf{P}(\mathfrak{L}^0_\Lambda = \{\mathbf{C}^0_1, \ldots, \mathbf{C}^0_{K^0}\},\, \mathfrak{L} = \{\mathbf{C}_1, \ldots, \mathbf{C}_K\}) = \frac{1}{Z^0_{\Lambda'}} \times \prod_{\ell=1}^{K^0} w(\mathcal{C}^0_\ell) \times \prod_{j=1}^{K} w(\mathcal{C}_j).$$

Summing over all instances of $\mathfrak{L}^0_\Lambda$, the independence of $\mathfrak{L}^0_\Lambda$ and $\mathfrak{L}$ implies

$$\mathbf{P}(\mathfrak{L} = \{\mathbf{C}_1, \ldots, \mathbf{C}_K\}) = \frac{1}{Z} \prod_{j=1}^{K} w(\mathcal{C}_j).$$

This proves the claim made in (10)

We introduce a partial order on elements of $\mathfrak{L}$ as follows: we say that $\mathbf{C} \prec \mathbf{C}'$, if there exist $j \geq 1$ and a sequence of coloured cycles $\mathbf{C} = \mathbf{C}_j, \mathbf{C}_{j-1}, \ldots, \mathbf{C}_0 = \mathbf{C}'$ all in $\mathfrak{L}$, such that for each

$1 \le r \le j$, the coloured cycles $\mathbf{C}_{r-1}$ and $\mathbf{C}_r$ share at least one vertex whose colour in $\mathbf{C}_r$ is one greater than its colour in $\mathbf{C}_{r-1}$. The meaning of the relation $\prec$ is the following:

$$\mathbf{C} \prec \mathbf{C}' \qquad \Longleftrightarrow \qquad \text{regardless of the order of popping, } \mathbf{C}' \text{ is popped before } \mathbf{C}. \tag{11}$$

(Recall that the set $\mathfrak{L}$ does not depend on the order of popping.) The direction $\Longrightarrow$ of this equivalence is immediate from the definition of $\prec$. To see the $\Longleftarrow$ direction, let us pop every cycle we can without popping $\mathbf{C}'$. This does not reveal $\mathbf{C}$. Now pop $\mathbf{C}'$, and note that any cycle that is revealed as a result of popping $\mathbf{C}'$ necessarily shares a vertex with $\mathbf{C}'$. Popping further cycles it holds that any cycle that is revealed has a chain of cycles leading to $\mathbf{C}'$. In particular, $\mathbf{C}$ must have this property. The equivalence (11) makes it clear that $\prec$ is a partial order on $\mathfrak{L}$.

We apply a parallel popping procedure to reveal $\mathfrak{L}$, defined in stages. In each stage, we pop all cycles on top of the stacks, simultaneously. If the event on the left hand side of (9) occurs, there exists a smallest integer $\ell \ge 1$, such that in stage ℓ we pop some cycle that intersects $\overline{W} := W_{k,\Lambda'} \cup \partial W_{k,\Lambda'}$. Indeed, if we never popped any such cycles, then the arrows attached to all the vertices in $\overline{W}$ would have the same direction as they had in T_Λ, which would force $W_{k,\Lambda'} = W_{k,\Lambda}$ and $T_{\Lambda'}|_{\mathsf{E}^*_{W,k}} = T_\Lambda|_{\mathsf{E}^*_{W,k}}$. Let us select, according to some fixed arbitrary rule, a cycle $\mathbf{D}_1 \in \mathfrak{L}$ popped in stage ℓ, and a vertex $w \in \mathbf{D}_1 \cap \overline{W}$. Let

$$\mathfrak{M} := \{\mathbf{D} \in \mathfrak{L} : \mathbf{D} \succeq \mathbf{D}_1\}. \tag{12}$$

Observe that $\mathfrak{M}$ can be popped from $\mathfrak{L}$ (without popping any other cycles), since by construction, $\mathfrak{M}$ is closed under domination in the partial order $\prec$. Define $\widetilde{\mathfrak{L}}$ to be the collection of coloured cycles left after popping $\mathfrak{M}$ from $\mathfrak{L}$.

Lemma 11. *The map $\mathfrak{L} \mapsto (\mathfrak{M}, \widetilde{\mathfrak{L}})$ is injective.* □

The proof is omitted as it immediately follows from the definition of the map.

We are going to join the cycles in $\mathfrak{M}$ into a single loop γ in $\mathbb{Z}^d$, and then bound the probability of the possible arising loops in Lemma 13 below. We set $\gamma(0) = w$. Note that by the definition of $\mathbf{D}_1$, the arrow at w is at the top of its stack. We define γ by following the arrows, starting with the one on the top of the stack of w, and whenever we visit a vertex v for the i-th time, we use the i-th coloured arrow at v. The walk stops upon the first return to w. We call γ the *loop associated to* $\mathfrak{M}$. The purpose of the next lemma is to show that γ is well-defined and the map $\mathfrak{M} \to \gamma$ is injective.

Lemma 12. *Let $W \subset D_n$ be a fixed set and let $w \in \partial W$ be a fixed vertex. Suppose that $\mathfrak{L}$ is a collection of coloured cycles that can be popped, and $\mathbf{D}_1 \in \mathfrak{L}$ has the property that $w \in \mathcal{D}_1$, but no coloured cycle popped at any earlier stage than $\mathbf{D}_1$ intersects $\overline{W} = W \cup \partial W$. Let $\mathfrak{M}$ be defined by formula* (12)*. Then we have:*
(i) The loop associated to $\mathfrak{M}$ is well-defined in that the walk does return to w.
(ii) Every coloured edge in $\mathfrak{M}$ is used exactly once by the loop.
(iii) The map $\mathfrak{M} \mapsto \gamma$ is injective.

Proof. (i), (ii) We prove the two statements together by induction on the number of cycles in $\mathfrak{M}$. If $\mathfrak{M}$ consists of the single cycle $\mathbf{D}_1$, the statement is trivial. Otherwise, consider the first time

we return to a vertex v that we visited before. Then the cycle just found, $\mathbf{D}$, say, is necessarily on top of the stacks and $\mathbf{D} \neq \mathbf{D}_1$. Also, since the walk starts with an arrow belonging to a cycle in $\mathfrak{M}$, it is easy to see that $\mathbf{D} \in \mathfrak{M}$. Now pop $\mathbf{D}$, and define $\mathfrak{L}'$, $\mathfrak{M}'$, $\mathbf{D}_1'$ by moving the arrows in the stacks of the vertices of $\mathcal{D}$ up by one (and removing the arrows in $\mathbf{D}_1$). Observe that $\mathfrak{L}'$, $\mathfrak{M}'$, $\mathbf{D}_1'$ also satisfy the hypotheses of the Lemma, so by the induction hypothesis, the walk γ' defined by $\mathfrak{M}'$ visits each arrow of $\mathfrak{M}'$ exactly once. Hence inserting into γ' the cycle $\mathcal{D}$ at v we get the walk γ defined by $\mathfrak{M}$. This implies the statements (i) and (ii).

(iii) This follows from the fact that by construction, following the history of the loop-erasure process on γ (started at w) the loops erased are precisely the loops in $\mathfrak{M}$. □

We continue with the proof of Proposition 10. We bound the left hand side of (9) from above as follows. Let Π denote the class of all sets of coloured loops that are possible values of $\mathfrak{L}$. Let Γ_w denote the collection of loops in $\mathbb{Z}^d$ that start and end at w and visit Λ^c. Let $\Gamma_{w,\Lambda'}$ denote those loops in Γ_w that stay inside Λ'. By the stated independence of the spanning tree in Λ' and $\mathfrak{L}$, we have

$$\begin{aligned}\mathbf{P}\left(W_{k,\Lambda'} \subset D_n \text{ and } W_{k,\Lambda'} \neq W_{k,\Lambda} \text{ or } T_{\Lambda'}|_{\mathsf{E}^*_{W,k}} \neq T_{\Lambda}|_{\mathsf{E}^*_{W,k}}\right) \\ \leq \sum_{W \subset D_n} \mu_{\Lambda'}(W_{k,\Lambda'} = W)\frac{1}{Z} \sum_{\substack{\mathfrak{L}\in\Pi:\exists \mathbf{D}_1\in\mathfrak{L},\\ \mathbf{D}_1\cap\partial W\neq\emptyset}} \prod_{\mathbf{C}\in\mathfrak{L}} w(\mathcal{C}).\end{aligned} \tag{13}$$

We fix W, and estimate the sum over $\mathfrak{L}$. To every $\mathfrak{L}$ occurring in the sum, we have associated (by our arbitrary rule), a choice of $w \in \partial W$ and $\mathfrak{M} \subset \mathfrak{L}$ containing w. This $\mathfrak{M}$, in turn determines a loop γ based at w. Observe that

$$\prod_{\mathbf{C}\in\mathfrak{L}} w(\mathcal{C}) = \prod_{\mathbf{D}\in\mathfrak{M}} w(\mathcal{D}) \times \prod_{\widetilde{\mathbf{C}}\in\widetilde{\mathfrak{L}}} w(\widetilde{\mathcal{C}}) = w(\gamma) \times \prod_{\widetilde{\mathbf{C}}\in\widetilde{\mathfrak{L}}} w(\widetilde{\mathcal{C}}).$$

Hence, using the injectivity statements in Lemma 11 and Lemma 12(iii), the right hand side of (13) is at most

$$\begin{aligned}\frac{1}{Z}\sum_{w\in\partial W}\sum_{\gamma\in\Gamma_{w,\Lambda'}} w(\gamma) \sum_{\widetilde{\mathfrak{L}}\in\Pi}\prod_{\widetilde{\mathbf{C}}\in\widetilde{\mathfrak{L}}} w(\widetilde{\mathcal{C}}) &\leq \sum_{w\in\partial W}\sum_{\gamma\in\Gamma_{w,\Lambda'}} w(\gamma) \\ &\leq \sum_{w\in\partial W}\sum_{\gamma\in\Gamma_w} w(\gamma).\end{aligned} \tag{14}$$

Lemma 13. *For any $w \in D_n$, we have*

$$\sum_{\gamma\in\Gamma_w} w(\gamma) \leq \frac{C_d}{(N-n)^{d-2}}. \tag{15}$$

Proof. The weight of a loop is equal to the probability of each step present occurring. Therefore the sum of the weights over loops Γ_w equals the sum of the probabilities of random walk paths

that start and end at w and exit Λ. Letting S denote simple random walk and τ_N the first exit time of D_N we get

$$\begin{aligned}\sum_{\gamma\in\Gamma_w} w(\gamma) &= \sum_{m\geq 0}\sum_{z\in\partial D_N}\sum_{r>m} \mathbf{P}^w(\tau_N = m,\, S(m)=z)\mathbf{P}^w(S(r)=w,\,|\,\tau_N=m,\,S(m)=z)\\ &= \sum_{m\geq 0}\sum_{z\in\partial D_N} \mathbf{P}^w(\tau_N=m,\,S(m)=z)G(z,w)\\ &\leq \frac{C_d}{(N-n)^{d-2}}\sum_{m\geq 0}\sum_{z\in\partial D_N}\mathbf{P}^w(\tau_N=m,\,S(m)=z)\\ &= \frac{C_d}{(N-n)^{d-2}}.\end{aligned}$$

Here $G(z,w)$ is Green's function, see [21, Section 4.3] for a proof of the bound on $G(z,w)$. □

Inserting (15) and (14) into (13) we get

$$\begin{aligned}&\mathbf{P}\left(W_{k,\Lambda'}\subset D_n \text{ and } W_{k,\Lambda'}\neq W_{k,\Lambda} \text{ or } T_{\Lambda'}|_{\mathsf{E}^*_{W,k}} \neq T_\Lambda|_{\mathsf{E}^*_{W,k}}\right)\\ &\quad\leq \frac{C_d}{(N-n)^{d-2}}\mathbf{E}_{\mu_{\Lambda'}}\left[|\partial W_{k,\Lambda'}| : W_{k,\Lambda'}\subset D_n\right]\\ &\quad\leq \frac{C_d}{(N-n)^{d-2}}\mathbf{E}_{\mu_{\Lambda'}}\left[|W_{k,\Lambda'}| : W_{k,\Lambda'}\subset D_n\right].\end{aligned} \tag{16}$$

We estimate the right hand side in the last equation in the following lemma.

Lemma 14. *We have*

$$\mathbf{E}_{\mu_{\Lambda'}}\left[|W_{k,\Lambda'}| : W_{k,\Lambda'}\subset D_n\right] \leq C_d k^{d-2}n^2.$$

Proof. By Wilson's algorithm, the probability that a vertex $x\in D_n\backslash D_k$ is in $W_{k,\Lambda'}$ is at most the probability that simple random walk started at x hits D_k. This is bounded by $C_d k^{d-2}/|x|^{d-2}$. Summing over $x\in D_n$ gives

$$\begin{aligned}\mathbf{E}_{\mu_{\Lambda'}}\left[|W_{k,\Lambda'}| : W_{k,\Lambda'}\subset D_n\right] &\leq |D_k| + \mathbf{E}_{\mu_{\Lambda'}}\left[|W_{k,\Lambda'}\cap(D_n\setminus D_k)|\right]\\ &\leq C_d k^d + C_d n^2 k^{d-2}\\ &\leq C_d n^2 k^{d-2}.\end{aligned}$$

□

The above lemma and (16) completes the proof of Proposition 10. □

Proposition 15. *Suppose $d\geq 3$. Then for sufficiently large n we have*

$$\mathbf{P}(W_k\not\subset D_n)\leq C_d k^{d-1}n^{\frac{2-d}{2d}}.$$

We prove this proposition by extending the argument of [22, Theorem 4.1], that requires a couple of alterations.

Proof. Condition on the event that the restriction of the uniform spanning forest to D_k, denoted $T|_{D_k}$, is a fixed forest K. Let K_j, $j = 1, 2, \dots$ denote the connected components of K. Then

$$\begin{aligned}\mathbf{P}(\mathrm{desc}(D_k) \not\subset D_n \,|\, T|_{D_k} = K) &= \mathbf{P}\big(\cup_j \big\{\mathrm{desc}(K_j) \not\subset D_n\big\} \,\big|\, T|_{D_k} = K\big) \\ &\leq \sum_j \mathbf{P}(\mathrm{desc}(K_j) \not\subset D_n \,|\, T|_{D_k} = K).\end{aligned}$$

In order to deal with the summand in the last expression, we need to generalize [22, Lemma 3.2]. Given a graph G, and V a subset of the vertices, we denote by G/V the graph obtained from G by identifying all the vertices in V to a single vertex and removing loop-edges.

Lemma 16. *Let G be a finite graph containing D_k as a subgraph and s a vertex of G with $s \notin D_k$. Let T_K denote the uniform spanning tree of G conditioned on its restriction to D_k being K. Let $L_j(T_K)$ denote the unique path from K_j to s in T_K. Then on the set of edges not belonging to K_j, the graph $T_K \setminus L_j(T_K)$ is stochastically dominated by the uniform spanning tree of $G/(K_j \cup \{s\})$, conditioned on the event that its restriction to D_k/K_j equals K/K_j.*

Proof. First we further condition on $L_j(T_K) = L$. Note that under this conditioning, $T_K \setminus L$ has the same distribution as the uniform spanning tree of $G/\mathrm{Vert}(L)$ given K, where $\mathrm{Vert}(\cdot)$ denotes vertex set of a graph. By the negative association theorem of Feder and Mihail [10], [23, Chapter 4], conditioning on an edge being present makes the remaining set of edges stochastically smaller. As $\mathrm{Vert}(L)$ contains both K_j and s we can repeatedly apply this result to deduce that on the edges not belonging to $K_j \cup L$ the set of edges $T_K \setminus L$ is dominated by the uniform spanning tree of $G/(K_j \cup \{s\})$ given K/K_j. We can now average over all possible paths $L = L_j(T_K)$ to remove this part of the conditioning and get the stated lemma. □

We will use the following corollary of Lemma 16 that can be deduced by taking weak limits. Let $\mathfrak{F}_{K,j}$ denote the wired spanning forest conditioned on K with K_j wired to infinity (defined as the weak limit of uniform spanning trees conditioned on K with K_j wired to the sink).

> The set of descendants of K_j in the wired uniform spanning forest conditioned on K is stochastically dominated by the connected component of K_j in $\mathfrak{F}_{K,j}$.

The rest of the proof follows an outline similar to the proof of [22, Theorem 4.1]. We define edge sets $E_1 \subset E_2 \subset \dots$ as follows. Let $E_0 = K_j$. Assuming E_n has been defined, let S_n be the set of vertices of the connected component of $\mathfrak{F}_{K,j} \cap E_n$ containing K_j. If all edges incident with S_n are in E_n, we set $E_{n+1} = E_n$. If not, let e be an edge incident with S_n that minimizes $\min\{r : e \subset B_r\}$, where $B_r = \{x \in \mathbb{Z}^d : \|x\|_\infty \leq r\}$, and set

$$E_{n+1} := \begin{cases} E_n \cup \{e\} & \text{if } e \text{ does not connect } S_n \text{ with a component } K_i,\ i \neq j; \\ E_n \cup \{e\} \cup K_i & \text{if } e \text{ connects } S_n \text{ with } K_i. \end{cases}$$

When in the above $E_n \subset B_{r-1}$, i.e. a "new shell is visited" by the process, we make the further requirement that e be the edge along which the unit current flow from S_n to ∞ is maximal.

Let M_n be the effective conductance from S_n to ∞ in the complement of E_n, with the edges of K shorted:

$$M_n := \mathcal{C}(S_n \leftrightarrow \infty \text{ in } (\mathbb{Z}^d/K) \setminus E_n).$$

Then by [22, Lemma 3.3], [27, Theorem 7], $(M_n)_{n\geq 0}$ is a martingale with respect to the filtration $\mathcal{F}_n$ generated by E_n, $\mathfrak{F}_{j,K} \cap E_n$.

The M_0 term is no longer constant, as in the original proof. Nevertheless, the argument of [22, Theorem 4.1] gives:

$$\mathbf{P}(\mathrm{desc}(K_j) \not\subset D_n \,|\, T|_{D_k} = K) \leq C_d n^{\frac{2-d}{2d}} M_0(K_j)$$

We now bound $M_0(K_j)$ still with the conditioning that on D_k we have the forest K. Therefore we can work on the graph produced by deleting any edges from D_k that do not appear in K and contracting each component of K to a distinct vertex. By definition, the effective conductance from K_j to ∞ is the infimum of the energy of functions that are zero on K_j and one except on finitely many vertices. Therefore consider the function defined by $g(v) = 0$ if $v \in K_j$ and one otherwise. This is clearly a valid function with regards to the infimum and will have energy equal to the number of edges connected to K_j. As all edges in D_k that are not present in K have been deleted and K_j is a connected component of K, the only edges will be those connected to K_j from the outside of D_k. The size of this set is at most $C_d|\partial D_k \cap K_j|$.

Summing over the connected components, and using the fact that the K_j's are disjoint and cover all of D_k, we get

$$\sum_j M_0(K_j) \leq C_d|\partial D_k| \leq C_d k^{d-1}.$$

Then as this bound is independent of K we can average over all possible K to get the unconditioned result:

$$\mathbf{P}(\mathrm{desc}(D_k) \not\subset D_n) \leq C_d n^{\frac{2-d}{2d}} k^{d-1}.$$

This completes the proof of Proposition 15. □

Proof of Theorem 9. If $W_k = W_{k,\Lambda}$ and T and T_Λ agree on $E^*_{W,k}$, then $\psi_{\mathcal{D}}$ and $\psi_{\mathcal{D},\Lambda}$ will produce the same sandpile configuration on D_k. Therefore to bound the difference of the measures on any cylinder event E defined on D_k it suffices to bound the probability that the descendants in the spanning trees differ, or the trees differ on that set of descendants.

$$\begin{aligned} |\nu(E) - \nu_\Lambda(E)| &\leq \mathbf{P}(E_{\mathbb{Z}^d} \Delta E_\Lambda) \\ &\leq \mathbf{P}\left(W_k \neq W_{k,\Lambda} \text{ or } T|_{\mathsf{E}^*_{W,k}} \neq T_\Lambda|_{\mathsf{E}^*_{W,k}}\right) \\ &\leq C_d \frac{k^{d-2}n^2}{(N-n)^{d-2}} + C_d k^{d-1} n^{\frac{2-d}{2d}}. \end{aligned} \tag{17}$$

We now optimise the choice of n. We may assume $N \geq 2n$, in which case $(N-n)^{d-2} \geq c_d N^{d-2}$.

When $d \geq 5$, we take $n = \frac{1}{2}N$, which gives the bound $C_d k^{d-1} N^{\frac{2-d}{2d}}$.

When $d = 4$, the two terms in the right hand side of (17) are of the same order if $n = k^{4/9}N^{8/9}$. This gives the bound $Ck^{26/9}N^{-2/9}$.

When $d = 3$, we take $n = k^{6/13}N^{6/13}$. This yields the bound $Ck^{17/9}N^{-1/9}$. □

4 Rate of convergence in $\mathbb{Z}^2$.

In this section we bound the rate of convergence on $\mathbb{Z}^2$ in Theorem 17 below. As was the case for $d \geq 3$, the result will follow directly from the bijections and a bound on the probability that, in a suitable coupling, the descendants of D_k in $\mathbb{Z}^2$ differ from those in Λ. This bound is given in Proposition 18. Due to recurrence, we cannot use Wilson's method rooted at infinity, so the construction of the coupling is more involved. Write $G = (\Lambda \cup \{s\}, \mathsf{E}_\Lambda)$ for the graph on which the sandpile is defined. Recall that given a cylinder event E determined by the sandpile heights in D_k, we write $E_{\mathbb{Z}^2} = \{\psi_{\mathcal{D}}(T) \in E\}$ and $E_\Lambda = \{\psi_{\mathcal{D},\Lambda}(T_\Lambda) \in E\}$, where T is a sample from WSF and T_Λ is a sample from UST_G.

Theorem 17. *Let E be a cylinder event determined by the sandpile heights in D_k, and let $\Lambda \subset \mathbb{Z}^2$ be a finite set. Let N be the largest integer such that $D_N \subset \Lambda$. Given $\varepsilon > 0$, there is a constant $C = C(\varepsilon) > 0$ and a coupling $\mathbf{P} = \mathbf{P}_{\Lambda,k,\varepsilon}$ of T and T_Λ, such that in this coupling we have*

$$|\nu(E) - \nu_\Lambda(E)| \leq \mathbf{P}(E_{\mathbb{Z}^2} \Delta E_\Lambda) \leq C \frac{k^{5/32}}{N^{1/16-\varepsilon}}.$$

We will write W_k, respectively $W_{k,\Lambda}$, for the set of descendants of D_k in T, respectively T_Λ. Then Theorem 17 follows immediately from the following proposition.

Proposition 18. *For any $\varepsilon > 0$ there exists $C = C(\varepsilon) > 0$ and a coupling $\mathbf{P} = \mathbf{P}_{\Lambda,k,\varepsilon}$ of T and T_Λ such that in this coupling*

$$\mathbf{P}\left(W_k \neq W_{k,\Lambda} \text{ or } T \text{ and } T_\Lambda \text{ differ on some edge touching } W_k\right) \leq C \frac{k^{5/32}}{N^{1/16-\varepsilon}}.$$

The coupling will be achieved by passing to the planar dual graphs. The idea is to construct paths in the dual tree that together surround D_k in such a way that all descendants of D_k are necessarily in the interior of the region defined by the paths. Then it will be sufficient to couple the dual trees in the interior of that region.

Let $G^* = (\Lambda^*, \mathsf{E}^*_\Lambda)$ denote the planar dual of G. The vertex set Λ^* is naturally identified with a subset of the dual lattice $(\mathbb{Z}^2)^* = \mathbb{Z}^2 + (1/2, 1/2)$. The planar graph G^* has one unbounded face: the face corresponding to the sink s via duality. The dual spanning tree T^*_Λ is defined on G^*, by including a dual edge e^* in T^*_Λ if and only if the corresponding edge e is not in T_Λ. Then T^*_Λ is a sample from UST_{G^*} (i.e. with free boundary conditions). It is well known that as $\Lambda \uparrow \mathbb{Z}^2$, the measure UST_{G^*} converges weakly to the free spanning forest measure FSF, which for $\mathbb{Z}^2$ coincides with WSF [28, 23]. Let T^* denote a sample from this measure on the graph $(\mathbb{Z}^2)^*$. We refer to paths in $\mathbb{Z}^2$ as *primal* paths, and paths in $(\mathbb{Z}^2)^*$ as *dual* paths. Let o^* be the dual vertex $o + (1/2, 1/2) \in (\mathbb{Z}^2)^*$, where o is the origin in $\mathbb{Z}^2$. For any $m \geq 0$ we define the balls in the dual graph:

$$D^*_m := \{w \in (\mathbb{Z}^2)^* : |w - o^*| \leq m\}.$$

The construction of the coupling is broken down into a sequence of steps, and the required estimates stated as lemmas. We collect the estimates at the end and prove Proposition 18. The integers $\ell \geq 1$ and $k < n < r < R < N$ will be parameters that we choose at the end to optimize the bound.

Step 1. Coupling the backbones inside D_r^*. We will need to work with fixed "backbones" in our trees. Since T^* has one end WSF-a.s., there is a unique infinite path γ^* in T^* that starts at o^*. We call γ^* the *backbone* of T^*. The free spanning tree on Λ^* does not have a unique backbone (there are typically several paths from o^* to the boundary of Λ^*). Therefore, we will first work with the *wired* boundary condition in the *dual graph*, i.e. we consider the graph $\widetilde{G}^* = (\Lambda^* \cup \{s^*\}, \widetilde{\mathsf{E}}_\Lambda^*)$ obtained by connecting each vertex in Λ^* to s^* by as many edges as it needs, for its degree to be 4. Then we will compare $\mathsf{UST}_{\widetilde{G}^*}$ to UST_{G^*} using the well known monotone coupling between them [28, 23]. Let $\widetilde{T}_\Lambda^*$ denote a sample from $\mathsf{UST}_{\widetilde{G}^*}$. Let γ_Λ^* denote the unique path between o^* and s^* in $\widetilde{T}_\Lambda^*$. We call γ_Λ^* the *backbone* of $\widetilde{T}_\Lambda^*$.

We fix a coupling between γ^* and γ_Λ^* that maximizes the probability that their first ℓ steps are identical. The next lemma collects some LERW estimates from the literature that we use to estimate the probability that the restrictions of γ_Λ^* and γ^* to the ball D_r^* differ from each other.

Lemma 19. *(i) For $l < \sqrt{N}$, we have*

$$\mathbf{P}(\textit{first } \ell \textit{ steps of } \gamma^* \textit{ and } \gamma_\Lambda^* \textit{ are not identical}) \leq C\frac{l^2}{N}\ln\left(\frac{N}{l}\right).$$

(ii) If $R > 4r$, we have

$$\mathbf{P}(\gamma_\Lambda^* \textit{ returns to } D_r^* \textit{ after its first exit from } D_R^*) \leq C\frac{r}{R}$$

and

$$\mathbf{P}(\gamma^* \textit{ returns to } D_r^* \textit{ after its first exit from } D_R^*) \leq C\frac{r}{R}.$$

(iii) We have

$$\mathbf{E}[\textit{number of steps of } \gamma^* \textit{ until first exit from } D_R^*] = R^{\frac{5}{4}+o(1)} \quad \textit{as } R \to \infty.$$

(iv) For all $\lambda, \varepsilon > 0$, $N > 4R$ we have that there exist $C(\varepsilon), C_1, C_2 > 0$ such that

$$\mathbf{P}(\textit{number of steps of } \gamma_\Lambda^* \textit{ until first exit from } D_R^* > \lambda C(\varepsilon) R^{\frac{5}{4}+\varepsilon}) \leq C_1 e^{-C_2\lambda}.$$

Remark 2. Note that in contrast with [21, Proposition 11.3.1], the above bounds give us power law (rather than logarithmic) control on the errors, since we are free to discard a collection of "bad paths" in D_r^* of small probability on which convergence to the infinite LERW would be much slower.

Proof of Lemma 19. (i) The statement follows from [20, Proposition 7.4.2]. Note that although the exact statement is not present in the reference, it immediately follows from the proof presented there.

(ii) This is [5, Lemma 2.4].

(iii) This result was first shown by Kenyon [19] (stated there in the upper half plane). It also follows by combining [4, Proposition 6.2(2)] and [26, Theorem 5.7].

(iv) This follows from [4, Corollary 3.4], [4, Theorem 5.8(4)] and part (iii). □

The next lemma puts the above estimates together and bounds the probability that the restrictions of γ^*_Λ and γ^* to the ball D^*_r are not identical.

Lemma 20.

$$\mathbf{P}(\gamma^*_\Lambda \cap D^*_r \neq \gamma^* \cap D^*_r) \leq C\frac{\lambda^2 C(\varepsilon)^2 R^{\frac{5}{2}+2\varepsilon}}{N} \ln\left(\frac{N}{\lambda C(\varepsilon) R^{5/4+\varepsilon}}\right) + C_1 \exp(-C_2\lambda) + 2C\frac{r}{R}$$

Proof. Let F_1 be the event that the first ℓ steps of γ^* and γ^*_Λ coincide, the event maximized by our choice of coupling. We therefore need to choose ℓ appropriately to get the desired result.

Let F_2 be the event that the length of γ^*_Λ until first exit of D^*_R is less than ℓ.

Let F_3 be the event that neither γ^*_Λ nor γ^* return to D^*_r after their first exits from D^*_R.

On the event $F_2 \cap F_3$, we have that the first ℓ steps of γ^*_Λ includes $\gamma^*_\Lambda \cap D^*_r$. If F_1 also occurs, then we have $\gamma^*_\Lambda \cap D^*_r = \gamma^* \cap D^*_r$. We choose $\ell = \lambda C(\varepsilon) R^{5/4+\varepsilon}$. By Lemma 19(i),(iv),(ii) we have

$$\begin{aligned}\mathbf{P}(\gamma^*_\Lambda \cap D^*_r \neq \gamma^* \cap D^*_r) &\leq \mathbf{P}(F_1^c) + \mathbf{P}(F_2^c) + \mathbf{P}(F_3^c) \\ &\leq C\frac{\lambda^2 C(\varepsilon)^2 R^{\frac{5}{2}+2\varepsilon}}{N} \ln\left(\frac{N}{\lambda C(\varepsilon) R^{5/4+\varepsilon}}\right) + C_1 \exp(-C_2\lambda) + 2C\frac{r}{R}.\end{aligned}$$

□

Step 2. Constructing the dual paths that surround D^*_k. On the event $\gamma^*_\Lambda \cap D^*_r \neq \gamma^* \cap D^*_r$, we extend the coupling of γ^*_Λ and γ^* to a coupling of WSF and $\mathsf{UST}_{\widetilde{G}^*_\Lambda}$ in an arbitrary way. (For example: make them conditionally independent given the backbones.) On the event $\gamma^*_\Lambda \cap D^*_r = \gamma^* \cap D^*_r$, we extend the coupling via Wilson's stacks of arrows construction. For each $x \in D^*_r \setminus \gamma^*$, we assign identical stacks for the constructions in Λ^* and $(\mathbb{Z}^2)^*$, respectively. For all other vertices, the stacks in Λ^* are assigned independently from those in $(\mathbb{Z}^2)^*$. This defines a coupling of WSF and $\mathsf{UST}_{\widetilde{G}^*}$ on $(\mathbb{Z}^2)^*$.

We now construct the required set of dual paths. Write γ^*_r for the portion of γ^* up to its first exit from D^*_r. By a *block*, we mean a set U of dual edges with the properties:
(i) $U \subset D^*_n \setminus D^*_k$;
(ii) $U \cup \gamma^*_r$ is a connected set of edges;
(iii) the set of vertices of $U \cup \gamma^*_r$ disconnects D^*_k from $(D^*_n)^c$.

Lemma 21. *Suppose that $r > 4n > 16k$ and $\gamma^*_\Lambda \cap D^*_r = \gamma^* \cap D^*_r$. There exists $C > 0$ such that*

$$\mathbf{P}\begin{pmatrix}\textit{we can pop a set of coloured cycles contained in} \\ D^*_n \setminus D^*_k \textit{ so that the arrows revealed contain a block}\end{pmatrix} \geq 1 - C\left(\frac{k}{n}\right)^{1/4} - C\frac{n}{r}.$$

Proof. Due to Lemma 19(ii), we have $\mathbf{P}(\gamma^* \cap D^*_n \neq \gamma^*_r \cap D^*_n) \leq C(n/r)$. Henceforth assume that we are on the event when $\gamma^* \cap D^*_n = \gamma^*_r \cap D^*_n$.

We start with a minor adaptation of the argument of [1, Lemma 6.1]. Let $v \in (\mathbb{Z}^2)^*$ be a vertex at distance $\sqrt{kn}$ from o^*, and let $\{S(n)\}_{n\geq 0}$ be simple random walk starting at v. Let τ be the first time when either S exits $D^*_n \setminus D^*_k$, or when the loop-erasure of S has made a non-contractible loop around D^*_k. Let us use the sequence $S(1), S(2), \ldots, S(\tau)$ as our successive

choices in Wilson's algorithm, where γ_r^* is already part of the tree to be constructed. That is, whenever a random walk step is to be made, we use the next step of S for the random walk step, and whenever a new vertex is to be chosen in the algorithm, we use the next vertex visited by S as the new vertex.

We claim that on the event $S[0,\tau] \subset D_n^* \setminus D_k^*$ the set of edges U that we have added to the tree is a block. Indeed, condition (i) holds because the walk never left $D_n^* \setminus D_k^*$. Also, observe that the set of vertices of $\mathrm{LE}(S[0,\tau))$ do not get erased, and hence condition (iii) holds. Finally, condition (ii) holds, because each piece of the tree we create gets joined to γ_r^* (here is where we use that $\gamma^* \cap D_n^* = \gamma_r^* \cap D_n^*$). Note that since $S(\tau-1)$ does not get erased, the last piece is also joined. This proves the claim. Interpreting the construction in terms of stacks of arrows, we see that the probability of the event in part (i) is at least the probability that $S[0,\tau] \subset D_n^* \setminus D_k^*$.

The probability that a non-contractible loop is created could be bounded by $\geq 1 - C(k/n)^\zeta$ with some $\zeta, C > 0$, by ideas similar to [21, Exercise 3.3], showing the statement (i) with ζ in place of $1/4$. In order to get the explicit exponent $1/4$, we combine the argument with an idea that was inspired by [6].

Again we are going to start with γ_r^* as our initial tree. Choose a subpath $\gamma_{k,n}^*$ of γ_r^* that forms a crossing from D_k^* to $(D_n^*)^c$. Write H_ρ for the circle of radius ρ centred at o^*. Define the annulus:

$$A_{k,n} = \{z \in \mathbb{R}^2 : k+1 < |z - o^*| < n-1\}.$$

Choose a point $Q \in \gamma_{k,n}^* \cap H_{\sqrt{kn}}$, and let $\alpha_0 = H_{\sqrt{kn}} \setminus \{Q\}$. Let P_1 be the mid-point of α_0, and let v_1 be a vertex of $(\mathbb{Z}^2)^*$ closest to P_1. Run a random walk $S^{(1)}$ from v_1 to $\gamma_{k,n}^*$, and add edges to the tree in the same way as we did with S. Let π_1 be the set of edges added. Note that π_1 is not necessarily a connected set of edges, however, $\gamma_r^* \cup \pi_1$ is. From the two subarcs of α_0 defined by P_1, throw away the one that is on the same side of $\gamma_{k,n}^*$ as where π_1 hit, and let us call the other arc α_1. On the event when $\{S^{(1)}\} \subset A_{k,n}$, the arc α_1 has the property that any dual lattice path from H_k to H_n that is vertex-disjoint from $\gamma_r^* \cup \pi_1$ has to intersect α_1.

Continue inductively in the following way. Suppose that for some $i \geq 1$ the arc α_i and the sets of edges $\pi_1, \ldots, \pi_i$ have been defined. Let P_{i+1} be the mid-point of α_i and let v_{i+1} be the vertex of $(\mathbb{Z}^2)^*$ closest to P_{i+1}. Run a random walk $S^{(i+1)}$ from v_{i+1} to $\gamma_{k,n}^* \cup \pi_1 \cup \cdots \cup \pi_i$, and let π_{i+1} be the set of edges that get added to the tree. From the two subarcs of α_i, throw away the one that is on the same side of γ_r^* as where π_{i+1} hit, and call the other one α_{i+1}. On the event when $\{S^{(i+1)}\} \subset A_{k,n}$, the arc α_{i+1} has the property that any dual lattice path from H_k to H_n that is vertex-disjoint from $\gamma_r^* \cup \pi_1 \cup \cdots \cup \pi_{i+1}$ has to intersect α_{i+1}.

The construction is well defined until a time when the length of the arc α_i becomes of order 1. Stop the construction the first time when $\mathrm{diam}(\alpha_i) < 10$, say. We can select further vertices $v_{i+1}, \ldots, v_{i+K}$ (with K a fixed constant, say, $K = \lceil 10\sqrt{2} + 4 \rceil$) such that if we start further random walks at these vertices, then $\gamma^* \cup \pi_1 \cup \cdots \cup \pi_{i+K}$ contains a block. An example of the start of this construction is shown in Figure 1.

It remains to bound the probability that the walks $S^{(1)}, S^{(2)}, \ldots$ all remain inside $D_n^* \setminus D_k^*$. The i-th walk $S^{(i)}$ starts at distance $O(2^{-i}\sqrt{kn})$ from the current tree $\mathfrak{T}_{i-1} := \gamma_{k,n}^* \cup \pi_1 \cup \cdots \cup \pi_{i-1}$. If it were to leave $D_n^* \setminus D_k^*$ without hitting $\mathfrak{T}_{i-1}$, it would first have to leave the ball

$$B^*(v_i; (1/4)\sqrt{kn}) := \{w \in (\mathbb{Z}^2)^* : |w - v_i| \leq (1/4)\sqrt{kn}\}.$$

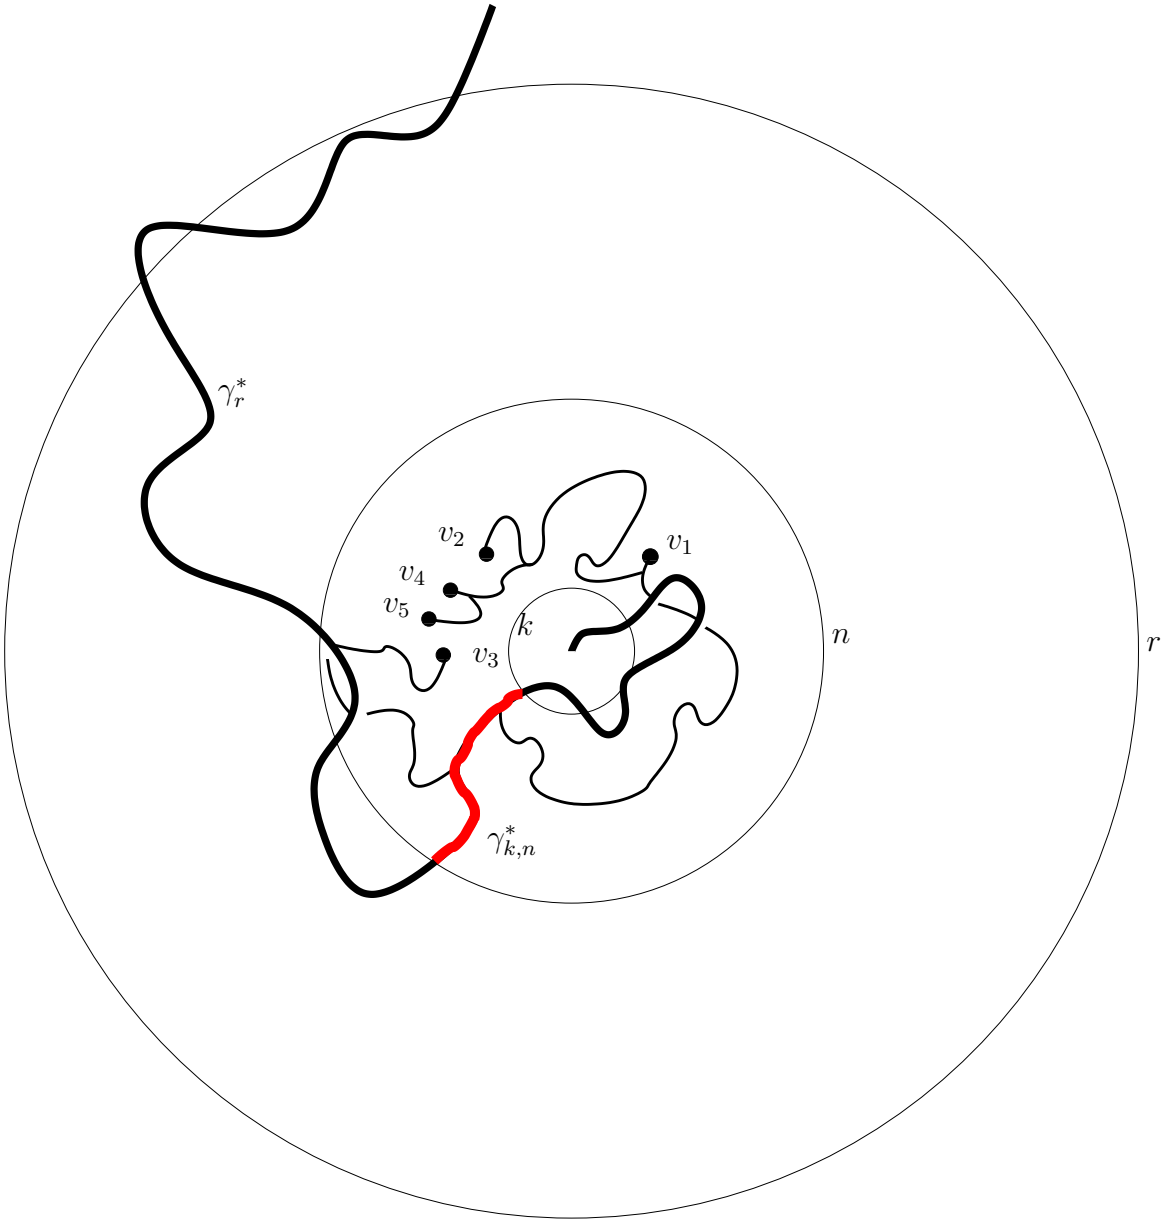

Figure 1: An example of the construction of a block. The thick line is γ_r^*, and the red piece is $\gamma_{k,n}^*$. LERWs were started successively at v_1, v_2, etc. Note the gaps between pieces in some of the LERWs, where an intersection with $\gamma_r^* \setminus \gamma_{k,n}^*$ has occurred.

without hitting $\mathfrak{T}_{i-1}$. Using Beurling's estimate [21, Section 6.8], the probability of this is at most $C(2^{-i}\sqrt{kn}/\sqrt{kn})^{1/2}$. Regardless of where the walk exits $B^*(v_i;(1/4)\sqrt{kn})$, the exit point z_i^* is still at distance $\asymp \sqrt{kn}$ from o^*. It follows, again using Beurling's estimate, that the probability that the continuation of the walk from z_i^* exits D_n^* without hitting $\mathfrak{T}_{i-1}$ is at most $C(\sqrt{kn}/n)^{1/2}$. Similarly, together with a time-reversal argument, the probability that the walk started at z_i^* hits D_k^* before hitting $\mathfrak{T}_{i-1}$ is at most $C(k/\sqrt{kn})^{1/2}$. Combining these three estimates we get the bound

$$\begin{aligned}\mathbf{P}(\{S^{(i)}\} \not\subset D_n^* \setminus D_k^* \,|\, \mathfrak{T}_{i-1}) &\le C\left(\frac{2^{-i}\sqrt{kn}}{\sqrt{kn}}\right)^{\frac{1}{2}} \times \left[\left(\frac{\sqrt{kn}}{n}\right)^{\frac{1}{2}} + \left(\frac{k}{\sqrt{kn}}\right)^{\frac{1}{2}}\right] \\ &= O\left(2^{-i/2}\left(\frac{k}{n}\right)^{\frac{1}{4}}\right).\end{aligned}$$

Summing over i we get the claimed bound $O((k/n)^{1/4})$. □

Step 3. Coupling the set of descendants. We now complete the definition of the coupling of T_Λ and T. Fix a monotone coupling between $\widetilde{T}_\Lambda^*$ and T_Λ^*, such that $e^* \in \widetilde{T}_\Lambda^* \cap \mathsf{E}_\Lambda^*$ implies $e^* \in T_\Lambda^*$; see [23, Chapter 10]. Define T_Λ and T as the dual trees of T_Λ^* and T^*. This completes the definition of required coupling $\mathbf{P}_{\Lambda,k,\varepsilon}$.

Lemma 22.
(i) When $\widetilde{T}_\Lambda^*$ contains a block, we have $W_{k,\Lambda} \subset D_r$.
(ii) When T^* contains a block, we have $W_k \subset D_r$.
(iii) When the event in Lemma 21 occurs, we have $W_{k,\Lambda} = W_k \subset D_r$ and T and T_Λ agree on the set of edges with at least one endvertex in W_k.

Proof. (i) Since T_Λ^* is stochastically larger than $\widetilde{T}_\Lambda^*$, the edges in the block are also present in T_Λ^*. Since the the union of the block with γ_r^* is connected, any two dual vertices in the interior of the block are connected by a path in $\widetilde{T}_\Lambda^*$. Hence no new edges are added in the interior of the block when passing from $\widetilde{T}_\Lambda^*$ to T_Λ^*.

Suppose that D_k had a descendant $v \in D_r^c$ in T_Λ. Then there would be a primal path β starting at v that visits D_k and ends outside D_r. Since the block surrounds D_k, this would contradict the connectivity of the block (as a set of edges).

(ii) The same argument as in the previous paragraph applies here.

(iii) Since we are using the same stacks of arrows in $D_r^* \setminus \gamma_r^*$, the same block exists in Λ^* and in $(\mathbb{Z}^2)^*$, and the trees coincide in the interior of the region defined by the block. Therefore, the trees T_Λ and T also coincide in this region. By parts (i) and (ii), the set of descendants are contained in this region and are equal in T_Λ and T. □

Proof of Proposition 18. By Lemma 22 we have $W_{k,\Lambda} = W_k$ if the event in Lemma 21 occurred which in turn assumed that the event in Lemma 20 did not occur. Therefore we have

$$\begin{aligned}&\mathbf{P}\left(W_k \neq W_{k,\Lambda} \text{ or } T \text{ and } T_\Lambda \text{ differ on some edge with an endvertex in } W_k\right) \\ &\quad\le C\frac{\lambda^2 C(\varepsilon)^2 R^{\frac{5}{2}+2\varepsilon}}{N}\ln\left(\frac{N}{\lambda C(\varepsilon)R^{5/4+\varepsilon}}\right) + C_1\exp(-C_2\lambda) + C\frac{r}{R} + C\left(\frac{k}{n}\right)^{1/4} + C\frac{n}{r}.\end{aligned}$$

We can now optimise our choice of parameters by taking $n = (kr^4)^{1/5}$, $r = (R^5k)^{1/6}$, $R = (kN^6)^{1/16}$ and choose λ such that $\lambda^2 R^{2\varepsilon} = N^{\varepsilon}$. □

Acknowledgements. We are grateful to an anonymous referee for useful suggestions.

References

[1] Aizenman, M.; Burcahrd, A.; Newman, C. M.; Wilson, D. B. Scaling limits for minimal and random spanning trees in two dimensions. Random Structures Algorithms **15** (1999), no. 3-4, 319–367.

[2] Athreya, S. R.; Járai, A. A. Infinite volume limit for the stationary distribution of abelian sandpile models. Comm. Math. Phys. **249** (2004), no. 1, 197–213.

[3] Bak, P. ; Tang, C.; Wiesenfeld, K. Self-organized criticality: An explanation of the 1/f noise. Physics Review Letters **59** (1987), 381–384

[4] Barlow, M. T.; Masson, R. Exponential tail bounds for loop-erased random walk in two dimensions. Ann. Probab. **38** (2010), no. 6, 2379–2417.

[5] Barlow, M. T.; Masson, R. Spectral dimension and random walks on the two dimensional uniform spanning tree. Comm. Math. Phys. **305** (2011), no. 1, 23–57.

[6] Benjamini, I. Large scale degrees and the number of spanning clusters for the uniform spanning tree, in *Perplexing problems in probability*, 175–183, Progr. Probab., 44, Birkhäuser, Boston, Boston, MA (2000).

[7] Benjamini, I.; Lyons, R.; Peres, Y. and Schramm, O. Uniform spanning forests. Ann. Probab. **29** (2001), no. 1, 1–65.

[8] Dhar, D. Self-organized critical state of sandpile automaton models. Phys. Rev. Lett. **64** (1990), no. 14, 1613–1616.

[9] Dhar, D. Theoretical studies of self-organized criticality. Phys. A **369** (2006), no. 1, 29–70.

[10] Feder, T. and Mihail, M. Balanced matroids. In Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing, pages 26–38, New York. Association for Computing Machinery (ACM). Held in Victoria, BC, Canada (1992).

[11] Holroyd, A. E.; Levine, L.; Mészáros, K.; Peres, Y.; Propp, J.; Wilson, D. B. Chip-firing and rotor-routing on directed graphs. In and out of equilibrium. 2, 331–364, Progr. Probab., **60**, Birkhäuser, Basel, (2008).

[12] Járai, A.A. Rate of convergence estimates for the zero dissipation limit in Abelian sandpiles. Preprint (2011),
`http://arxiv.org/abs/1101.1437`.

[13] Járai, A.A. Sandpile models. Preprint (2014),
`http://arxiv.org/abs/1401.0354`.

[14] Járai, A.A. and Lyons, R. Ladder sandpiles. Markov Process. Related Fields **13** (2007), no. 3, 493–518.

[15] Járai, A.A.; Werning, N. Minimal configurations and sandpile measures. J. Theoret. Probab. **27** (2014), no. 1, 153–167.

[16] Járai, A.A. and Redig, F. Infinite volume limit of the abelian sandpile model in dimensions $d \geq 3$. Probab. Theory Related Fields **141** (2008), no. 1–2, 181–212.

[17] Járai, A.A.; Redig, F. and Saada, E. Approaching criticality via the zero dissipation limit in the abelian avalanche model. Preprint (2014), `http://arxiv.org/abs/0906.3128`.

[18] Jeng, M.; Piroux, G. and Ruelle, P. Height variables in the Abelian sandpile model: scaling fields and correlations. J. Stat. Mech. Theory Exp. (2006), no. 10, P10015.

[19] Kenyon, R. The asymptotic determinant of the discrete Laplacian. Acta Math. **185** (2000), no. 2, 239–286.

[20] Lawler, G. F. Intersections of random walks. Reprint of the 1996 edition. Modern Birkhäuser Classics. Birkhäuser/Springer, New York, (2013).

[21] Lawler, G. F.; Limic, V. Random walk: a modern introduction. Cambridge Studies in Advanced Mathematics **123**. Cambridge University Press, Cambridge, (2010).

[22] Lyons, R.; Morris, B. J.; Schramm, O. Ends in uniform spanning forests. Electron. J. Probab. **13** (2008), no. 58, 1702–1725.

[23] Lyons, R. with Peres, Y. *Probability on Trees and Networks*. Cambridge University Press. In preparation (2014). Current version available at `http://mypage.iu.edu/~rdlyons/`.

[24] Majumdar, S.N. and Dhar, D. Height correlations in the Abelian sandpile model. J. Phys. A **24** (1991), no. 7, L357–L362.

[25] Majumdar, S.N.; Dhar, D. Equivalence between the Abelian sandpile model and the $q \to 0$ limit of the Potts model. J. Phys. A **185** (1991), 129–145.

[26] Masson, R. The growth exponent for planar loop-erased random walk. Electron. J. Probab. **14** (2009), no. 36, 10121073.

[27] Morris, B.J. The components of the wired spanning forest are recurrent. Probab. Theory Related Fields **125** (2003), 259–265.

[28] Pemantle, R. Choosing a spanning tree for the integer lattice uniformly. Ann. Probab. **19** (1991), no. 4, 1559–1574.

[29] Priezzhev, V.B. Structure of two-dimensional sandpile. I. Height probabilities. J. Stat. Phys. **74** (1994), no. 5–6, 955–979.

[30] Redig, F. Mathematical aspects of the abelian sandpile model, in *Mathematical statistical physics*, 657–729, Elsevier B. V., Amsterdam, (2006).

[31] Schmidt, K. and Verbitskiy, E. Abelian sandpiles and the harmonic model. Comm. Math. Phys. **292** (2009), no. 3, 721–759.

[32] Wilson, D.B. (1996): Generating random spanning trees more quickly than the cover time. Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996), 296–303, ACM, New York.